Service Chain (SC) Mapping with Multiple SC Instances in a Wide Area Network

This is a preprint electronic version of the article submitted to IEEE GlobeCom 2017

Abhishek Gupta*, Brigitte Jaumard†, Massimo Tornatore*‡, and Biswanath Mukherjee*
*University of California, Davis, USA †Concordia University, Canada ‡Politecnico di Milano, Italy
Email: *{abgupta,mtornatore,bmukherjee}@ucdavis.edu †bjaumard@cse.concordia.ca ‡massimo.tornatore@polimi.it

arXiv:1704.06716v1 [cs.NI] 21 Apr 2017

***Abstract*—Network Function Virtualization (NFV) aims to simplify deployment of network services by running Virtual Network Functions (VNFs) on commercial off-the-shelf servers. Service deployment involves placement of VNFs and in-sequence routing of traffic flows through VNFs comprising a Service Chain (SC). The joint VNF placement and traffic routing is usually referred as SC mapping. In a Wide Area Network (WAN), a situation may arise where several traffic flows, generated by many distributed node pairs, require the same SC, one single instance (or occurrence) of that SC might not be enough. SC mapping with multiple SC instances for the same SC turns out to be a very complex problem, since the sequential traversal of VNFs has to be maintained while accounting for traffic flows in various directions.**

Our study is the first to deal with SC mapping with multiple SC instances to minimize network resource consumption. Exact mathematical modeling of this problem results in a quadratic formulation. We propose a two-phase column-generation-based model and solution in order to get results over large network topologies within reasonable computational times. Using such an approach, we observe that an appropriate choice of only a small set of SC instances can lead to solution very close to the minimum bandwidth consumption.

I. Introduction

Today's communication networks deploy network services through proprietary hardware appliances (e.g., network functions such as firewalls, NAT, etc.) which are statically configured. With rapid evolution of applications, however, networks require more agile and scalable service deployment.

Network Function Virtualization (NFV) [1] offers a solution for more agile service deployment. NFV envisions hardware functionality as software modules called Virtual Network Functions (VNFs). VNFs can be run on commercial-off-the-shelf hardware such as servers and switches in datacenters (DCs), making service deployment agile and scalable.

When several network functions are configured to provide a service, we have a "Service Chain". The term "service chain" is used "to describe the deployment of such functions, and the network operator's process of specifying an ordered list of service functions that should be applied to a deterministic set of traffic flows" [2]. So, a "Service Chain" (SC) specifies a set of network functions configured in a specific order. With NFV, we can form SCs where VNFs are configured in a specific sequence that allows the minimization of the bandwidth usage in the network (an example is discussed in the next paragaph).

Unfortunately, the fact that VNFs in a single SC need to be traversed by several distinct traffic flows (i.e., flows requested by multiple geographically-distributed node pairs) in a specific sequence makes it difficult to improve network resource utilization. As an example, let us refer to Figs. 1(a) and 1(b), where three traffic requests r_1 (from node 4 to 13), r_2 (from node 6 to 3), and, r_3 (from node 14 to 1) demand SC c_1 composed of VNF1, VNF2 and VNF3 (to be traversed in this order VNF1, VNF2 and VNF3). In Fig. 1(a), we see that if we consider only one mapping occurrence for SC c_1, then some traffic flows (in our example, r_3 and r_2) will be ineffectively routed over a very long path. Instead, as shown in Fig. 1(b), if we use two SC mappings for the same SC, we can significantly improve network resource utilization, at the expense of a larger number of VNFs to be deployed/replicated in the network to serve the same SC. It results in an even more complex problem when in a Wide Area Network (WAN) a very large number of distributed node pairs generate traffic flow, creating a heavily-populated (dense) traffic matrix. Our objective in this work is to reduce the network resource consumption for a WAN with a dense traffic matrix.

So the question that arises is: how many SC instances for the same SC are required to achieve a quasi-optimal network resource utilization?

A possible trivial solution to the problem of SC mapping in case of multiple node pairs requiring the same SC is to use one single instance, that would most likely lead to host SCs at a single node (e.g., a DC) which is centrally located in the network. However, traffic flows will have to take long paths in order to reach the node hosting the SC, which results in a very suboptimal network resource consumption.

The other extreme case would be to use a distinct SC mapping per each node pair (in other words, the number of SC instances is equal to the number of traffic node pairs). In such a way, we can achieve optimal network resource utilization as each node pair will use an SC effectively mapped along a shortest path in the network. However, this approach has the downside of increasing the network orchestration overhead and increase capital expenditure, as there will be a very large number of replicated VNF instances across nodes. To reduce excessive VNF replication, we bound the maximum number of nodes hosting VNFs. Using the shortest path also has the added effect of reducing latency for the service chain, but this aspect is out of scope for this study.

Intuitively, the number of SC instances for a good solution will be a value between these two extreme points. This

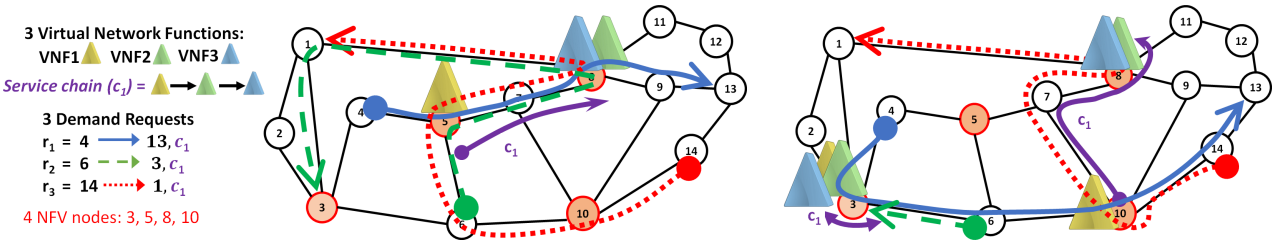

(a) One single SC instance mapping: request r_3 and r_2 take longer paths

(b) Two SC instances mapping: request r_3 and r_2 take shorter paths

Fig. 1: Deploying more SC occurrence mappings reduces network resource consumption.

solution will minimize the network resource utilization while not excessively increasing the number of nodes hosting VNFs.

A reasonable trade-off value is difficult to be optimally calculated. In fact, it has been shown that the problem of SC mapping with multiple SC instances results in quadratic constraints [3], that severely hamper the scalability of the solution. In this paper, to provide an answer to the question above, we propose a two-phase solution, relying on an ILP column-generation-based model, which provides quasi-optimal solutions with reasonable computational times. Sub-optimality comes from the fact that we solve the problem in two phases: in the first phase we group node pairs that will be forced to use the same SC instance, in the second phase we run our scalable column-generation approach to find a solution starting from the grouping already performed in the first phase. Applying the proposed approach over two realistic network topologies, we observe that an appropriate choice of only a small set of different SC mappings can lead to a solution very close to the minimum theoretical bandwidth consumption, even for a full-mesh traffic matrix.

The rest of this study is organized as follows. Section II overviews the existing literature on the SC mapping problem and remarks the novel contributions of this study. Section III formally describes the problem and its input parameters. Section IV describes a heuristic to cluster groups of node pairs that will use the same SC instance. We then describe our column-generation-based solving method in Section V. Section VI provides some illustrative examples that demonstrate that a limited number of SC instances can lead to quasi-optimal solution of the problem. Section VII concludes the study.

II. Related Work

A number of studies exist on the VNF placement and routing problem. Mehraghdam *et al.* [3] were the first to formally define the problem of VNF placement and routing. However, they developed a Quadratic Constrained Program (QCP), making it unscalable beyond small problem instances. The authors of Ref. [4] study a hybrid deployment scenario with hardware middleboxes using an Integer Linear Program (ILP), but do not enforce VNF service chaining explicitly. Ref. [5] uses an ILP to study trade-offs between legacy and NFV-based traffic engineering but does not have explicit VNF service chaining. Ref. [6] models the problem in a DC setting using an ILP to reduce the end-to-end delays and minimize resource over-provisioning while providing a heuristic to do the same. Here too the VNF service chaining is not explicitly enforced by the model. Ref. [7] models the batch deployment of multiple chains using an ILP and develops heuristics to solve larger instances of the problem. However, they enforce that VNF instances of a function need to be on a single machine and restrict all chains to three VNFs. Our model does not impose such constraints, and we allow any VNF type to be placed on any node and any number of VNFs in a SC while service chaining VNFs for a SC explicitly. Ref. [8] accounts for the explicit service chaining of VNFs but focuses on compute resource sharing among VNFs. Ref. [9] also uses a column generation model to solve the VNF placement and routing but considers dedicated SC instances per each traffic pair, hence solving the second extreme case mentioned in the introduction, which is a particular case of our approach.

Our previous work [10] and most of existing works solve the problem for multiple SCs but a single instance of the SC. We remark again that in the current work we consider multiple SCs, but with multiple instances per SC, hence most of the existing works represent a particular case of our current work where, each node pair requesting an SC has its own instance. Further, we also consider multiple geographically distributed node pairs which make a heavily-populated (dense) traffic matrix. Unfortunately, scaling the model to multiple instances per SC results in quadratic constraints. Hence, we now propose a novel decomposition model (column generation) for SC mapping with multiple SC instances (Section IV), that, together with a traffic-grouping heuristic, allows the solution of the problem for multiple same-SC requests in the same SC instance (Section V). Our objective is to minimize network-resource consumption while keeping a bound on the number of nodes that can host VNFs.

To the best of our knowledge, this is the first attempt to address the solution of the complete SC mapping problem (i.e., with multiple SC instances) over large network instances.

III. Problem Description

An operator's network provides multiple services and each service is realized by traversing a Service Chain (SC). Here, we assume that, for each service, the operator knows the ordered sequence of VNFs forming SC. To provide multiple

services, a network operator has to map corresponding SCs into network. This is the problem of multiple SC mapping. Most recent works on this topic assume SC mapping with a single SC instance for each SC i.e., all demands for a SC will be mapped to a single instance. Our work takes this further by allowing SC mapping to multiple SC instances, for each SC, which becomes a more complex problem. We solve this problem by using a two phases approach that will be explained in this the following sections.

A. Problem Statement

Given a network topology, capacity of links, a set of network nodes with NFV support (NFV nodes), compute resources at NFV nodes, the maximum number of NFV-nodes that can be used, traffic flows for source-destination pairs requiring a specific SC with a certain bandwidth demand, set of VNFs, and, set of SCs, we determine the placement of VNFs and corresponding traffic routing to minimize network-resource (bandwidth) consumption. Note that VNFs can be shared among different SCs.

B. Input Parameters

- $G = (V, L)$: Physical topology of backbone network; V is set of nodes and L is set of links.
- $V^{\text{NFV}} \subseteq V$: Set of NFV nodes.
- K: number of NFV nodes allowed to host VNFs.
- F, indexed by f: Set of VNFs.
- n^{CORE}: Number of CPU cores present per NFV node.
- n_f^{CORE}: Number of CPU cores per Gbps for function f.
- C: Set of chains, indexed by c.
- n_c: Number of VNFs in SC c.
- $\mathcal{SD}$: Set of source-destination (v_s, v_d) pairs.
- D_{sd}^c: Traffic demand between v_s and v_d for SC c.
- $\sigma_i(c)$: ID of ith VNF in SC c where $f_{\sigma_i(c)} \in F$.

As already mentioned, we are solving this problem considering that each SC request can map to multiple instances, which makes the problem quadratic. To avoid quadratic constraints, our solution approach has two phases:

- Phase 1: We fix the number N_c of instances accepted per SC (N_c can go from 1 up to the number of demands for that SC), and then we group the traffic requests in N_c groups of requests. All the requests in a group are forced to use the same SC instance (Section IV). Then we pass the N_c instances as distinct SCs to the next phase.
- Phase 2: We solve the SC mapping problem with one single instance per SC based on the the inputs of Phase 1. The solution of this simplified (linear, yet still very complex) problem is based on a column-generation-based decomposition model (Section V).

IV. Phase 1: Shortest Path Traffic Grouping (SPTG) Heuristic

In this section, we propose a Shortest Path Traffic Grouping (SPTG) heuristic, which forms N_c groups of node pairs for each SC (given by SD_c), to be given as input to the decomposition model in Section V that will treat them as distinct SC and decide the best SC mapping for each of the N_c node-pair groups. As a result, we will have a solution mapping multiple SC instances per SC.

The logic of the SPTG algorithm below is that groups are formed among node pairs that share links along their shortest path(s). SPTG uses the links on the shortest paths of node pairs to group node pairs together.

Algorithm 1 SPTG(c)

Require: G, SD_c, N_c

Ensure: PARTITION $\leftarrow$ partition of node pairs (v_s, v_d) into groups

PARTITION $\leftarrow \emptyset$
$numberOfGroups \leftarrow 0$
$SD_c^{\text{LEFT}} \leftarrow SD_c$ ▷ list of (v_s, v_d) for c
while $numberOfGroups < N_c$ & $SD_c^{\text{LEFT}} \neq \emptyset$ **do**
  for (v_s, v_d) in G **do**
    CLUSTER$_{sd}$ $\leftarrow$ set of traffic pairs whose shortest path passes through (v_s, v_d)
  end for
  $largestCluster \leftarrow \max_{(v_s,v_d):D_{sd}^c>0}$ CLUSTER$_{sd}$
  $SD_c^{\text{LEFT}} \leftarrow SD_c^{\text{LEFT}} \setminus largestCluster$ ▷ remove traffic pairs of $largestCluster$ from SD_c^{LEFT}
  Add $largestCluster$ to PARTITION
  $numberOfGroups \leftarrow numberOfGroups + 1$
end while
if $SD_c^{\text{LEFT}} \neq \emptyset$ **then**
  for $trafficPair \in SD_c^{\text{LEFT}}$ **do**
    add $trafficPair$ to GROUP $\in$ PARTITION, such that the (v_s, v_d) associated with GROUP provides the shortest path for provisioning $trafficPair$
  end for
end if

If Algorithm 1 terminates with $SD_c^{\text{LEFT}} = \emptyset$ and a number of groups that is $< N_c$, partition some of the groups in order to reach N_c groups.

V. Phase 2: Column-Generation Approach

To present the decomposition a.k.a column generation model, we use the following notation for SC representation. Each SC, denoted by c, is characterized by an ordered set of n_c functions:

$$[\text{SC } c] \qquad f_{\sigma_1(c)} \prec f_{\sigma_2(c)} \prec \cdots \prec f_{\sigma_{n_c}(c)} \tag{1}$$

Each deployment of SC c is defined by a set of VNF locations, and a set of paths, from the location of first VNF to location of last VNF.

Our decomposition model is based on a set of *SC configurations* where each configuration is associated with a potential provisioning of a SC c and a potential node placement of its functions. Let Γ be the set of configurations, and Γ_c be the subset of configurations associated with service chain (note that, to keep the problem linear, only one configuration per SC must be selected) $c \in C$: $\Gamma = \bigcup_{c \in C} \Gamma_c$.

As the number of potential configurations grows exponentially with network size, we find the problem to fit naturally in the column-generation framework [11].

Column generation (**CG**) is a decomposition technique, where the problem (called Master Problem - **MP**) to be solved is divided into two sub-problems: restricted master problem (**RMP**) (selection of the best configurations) and pricing problems (**PP_SC**(c))$_{c\in C}$ (configuration generators for each chain). The CG process involves solving the **RMP**, querying the dual values of **RMP** constraints, and using them for **PP_SC**(c) objective. Each improving solution (i.e., with a negative reduced cost) of **PP_SC**(c) is added to **RMP**, and previous step is repeated until optimality condition is reached (see [11], [12]), with the **PP_SC**(c) explored in a round robin fashion. A chain configuration is characterized by the following parameters:

- Location of the functions: $a_{vi}^{\gamma} = 1$ if ith function $f_i \in c$ is located in v in configuration; 0 otherwise.
- Connectivity of the locations: path from the location of current VNF to next VNF in SC c. If link ℓ is used in the path from the location of $f_{\sigma_i(c)}$ to the location of $f_{\sigma_{i+1}(c)}$, then $b_{i\ell}^{\gamma} = 1$; 0 otherwise.

A. *Restricted Master Problem (**RMP**)*

RMP selects the best $\gamma \in \Gamma_c$ for each SC c. Also it finds a route from v_s (source) to first VNF of c and from last VNF of c to v_d (destination).

An illustration of the constraint splitting between **RMP** and **PP_SC**(c) is depicted in Fig. 2. Nodes circled in purple are NFV nodes, yellow nodes do not host VNFs at present but have NFV support, and orange nodes currently host VNFs. Figure 2(a) has f_1 located at v_1. When a different configuration is selected in Fig. 2(b) and f_1 is located at v_2, then **RMP** finds the path from v_s to location of f_1. Similarly, **RMP** finds the path from last VNF to v_d, i.e., f_5 to v_d here.

Variables:

- $z_\gamma = 1$ if configuration γ is selected; 0 otherwise.
- $x_v^{ci} = 1$ if ith function of c is located in v; 0 otherwise.
- $y_\ell^{\text{first}(c),sd} = 1$ if ℓ is on path from v_s to location of first VNF in c; 0 otherwise.
- $y_\ell^{\text{last}(c),sd} = 1$ if ℓ is on path from location of last VNF in c to v_d; 0 otherwise.
- $h_v = 1$ if v is used as a location for a VNF; 0 otherwise.

Objective: Minimize bandwidth consumed:

$$\min \sum_{\gamma\in\Gamma} \overbrace{\underbrace{\left(\sum_{(s,d)\in\mathcal{SD}} D_{sd}^c\right)}_{\text{Overall traffic using } c}\underbrace{\left(\sum_{\ell\in L}\sum_{i\in I} b_{i\ell}^{\gamma}\right)}_{\text{Number of links in the route of } c} z_\gamma}^{\text{COST}_\gamma} + \sum_{c\in C}\sum_{\ell\in L}\sum_{(s,d)\in\mathcal{SD}} D_{sd}^c \left(y_\ell^{f_1(c),sd} + y_\ell^{f_{n_c}(c),sd}\right). \quad (2)$$

Total bandwidth consumed in placing multiple SCs depends on configuration γ selected for each SC c. Each γ for c locates VNFs of c and gives the route to traverse these VNF locations. So, bandwidth consumed when going from v_s to v_d and traversing the SC depends on selected γ.

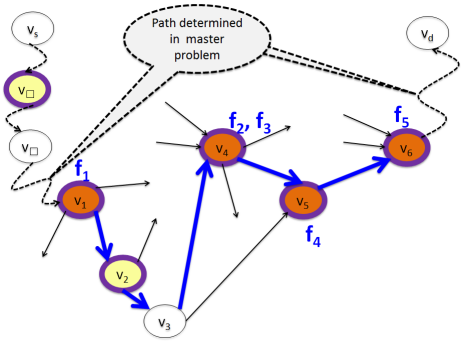

(a) A first configuration (γ_1) for c

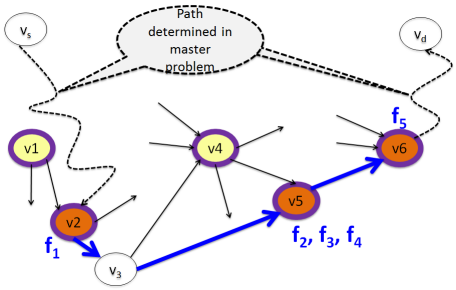

(b) A second configuration (γ_2) for c

Fig. 2: Two configuration examples for chain $c = (f_1 \prec f_2 \prec f_3 \prec f_4 \prec f_5)$.

Constraints:

$$\sum_{\gamma\in\Gamma_c} z_\gamma = 1 \qquad c \in C \quad (3)$$

$$\sum_{c\in C}\sum_{\gamma\in\Gamma_c}\sum_{(v_s,v_d)\in\mathcal{SD}} D_{sd}^c \left(\sum_{i=1}^{n_c} a_{vi}^{\gamma}\delta_{fi}^c n_f^{\text{CORE}}\right) z_\gamma \le n^{\text{CORE}} \qquad v \in V^{\text{NFV}} \quad (4)$$

$$\sum_{c\in C}\sum_{(v_s,v_d)\in\mathcal{SD}} D_{sd}^c \left(y_\ell^{f_1(c),sd} + y_\ell^{f_{n_c}(c),sd} + \sum_{\gamma\in\Gamma_c}\sum_{i=1}^{n_c-1} b_{i\ell}^{\gamma} z_\gamma\right) \le \text{CAP}_\ell \qquad \ell \in L \quad (5)$$

$$\sum_{\gamma\in\Gamma_c} a_{vi}^{\gamma} z_\gamma = x_v^{ci} \qquad f_i \in F(c), c \in C, v \in V^{\text{NFV}} \quad (6)$$

$$Mx_{vf} \ge \sum_{c\in C: f\in c}\ \sum_{i\in\{1,2,\ldots,n_c\}: f_i = f} x_v^{ci} \ge x_{vf} \qquad v \in V^{\text{NFV}}, f_i \in F \quad (7)$$

$$Mh_v \ge \sum_{f\in F} x_{vf} \ge h_v \qquad v \in V^{\text{NFV}} \quad (8)$$

$$\sum_{v\in V^{\text{NFV}}} h_v \le K. \quad (9)$$

Constraints (3) guarantee that we select exactly one γ for SC c and forces c to have a single instance. Each γ is associated with a set of a_{vi}^{γ} (from **PP_SC(c)**) required to be consistent with x_v^{ci} in **RMP**, which is resolved by Eqs. (6).

Constraints (4) ensure that each NFV node has a sufficient number of CPU cores for hosting f. Those constraints also accounts for increase in compute resource due to traffic increase. Eq. (5) enforces link-capacity constraints for the complete route for SC c from v_s to v_d for all $(v_s, v_d) \in \mathcal{SD} : D_{sd}^c > 0$).

Eq. (7) is used to keep track of VNF replicas. Eq. (8) is used to keep track of NFV nodes used for hosting VNFs while Eq. (9) enforces the number of NFV nodes allowed to host VNFs.

Route from v_s to first function location:

$$\sum_{\ell \in \omega^+(v_s)} y_\ell^{f_1(c),sd} = 1 - x_{v_s}^{c,1} \qquad c \in C, (v_s, v_d) \in \mathcal{SD} : D_{sd}^c > 0 \tag{10}$$

$$\sum_{\ell \in \omega^-(v)} y_\ell^{f_1(c),sd} \geq x_v^{c,1} \qquad c \in C, (v_s, v_d) \in \mathcal{SD} : D_{sd}^c > 0, v \in V^{\text{NFV}} \setminus \{v_s\} \tag{11}$$

$$\sum_{\ell \in \omega^+(v)} y_\ell^{f_1(c),sd} - \sum_{\ell \in \omega^-(v)} y_\ell^{f_1(c),sd} = -x_v^{c,1}$$
$$c \in C, (v_s, v_d) \in \mathcal{SD} : D_{sd}^c > 0, v \in V^{\text{NFV}} \setminus \{v_s\} \tag{12}$$

$$\sum_{\ell \in \omega^+(v)} y_\ell^{f_1(c),sd} - \sum_{\ell \in \omega^-(v)} y_\ell^{f_1(c),sd} = 0$$
$$c \in C, (v_s, v_d) \in \mathcal{SD} : D_{sd}^c > 0, v \in V \setminus (V^{\text{NFV}} \cup \{v_s\}). \tag{13}$$

We assume that a unique route exists from v_s to first VNF location. This is imposed by selecting exactly one outgoing link from v_s unless first VNF is located at v_s. We account for these scenarios using Eq. (10). To find the route from v_s to first VNF, flow conservation needs to be enforced at the intermediate nodes which may or may not have NFV support. Eqs. (12) and (13) enforces flow-conservation constraints at nodes with and without NFV support, respectively.

We can enforce same functionality as Eqs. (10), (12), (13), and (11), on route from location of last VNF to v_d. For the interested reader, similar details are provided in [13].

B. Pricing Problem

Mapping computations for each SC c ($c \in C$) correspond to the solution of pricing problems. The number of pricing problems to be solved equals the sum of the number of SC instances to be deployed.

Pricing problem **PP_SC(c)** generates: *(i)* A set of locations for VNFs of c; and *(ii)* a sequence of paths from the location of VNF f_i to the location of VNF f_{i+1}, for $i = 1, 2, \ldots, n_c - 1$ for chain c. Each solution that is generated by **PP_SC(c)** with a negative reduced cost, leads to a new potential γ for c of interest. Please refer to [13] for similar details.

C. Solution Scheme

The **PP_SC(c)** are solved in a round-robin fashion and the final **RMP** is solved as an ILP, as in [11], [12].

VI. Illustrative Numerical Examples

We first tested our two-phase optimization process on a 14 node NSFNet WAN topology [13] with a complete traffic matrix, i.e., with traffic flows between all node pairs, assuming all nodes can be made NFV nodes. The link capacity is sufficient to support all flows. Each traffic flow is 1 Gbps and demands the same 3 VNF service chain (SC). Compute resource (CPU) at each node is sufficient to support all VNF placements. The second run of the model is on a 11 node COST239 WAN topology [14] under the same specifications as above.

Figure 3(a) shows the bandwidth consumption as the number of SC instances increases. Here, we allow all nodes (K=14) to host VNFs. We find that as the number of deployed SC instances increases, the bandwidth consumption decreases. Indeed, with a higher number of instances, groups of traffic node pairs are able to take short paths. We see that at 34 instances, we achieve minimum possible bandwidth consumption, meaning traffic flow is taking the shortest paths. Note that the number of traffic node pairs in the network is 182, requiring a priori 182 different instances (solving the problem for 182 instances would be equivalent to obtaining a solution with existing models as in Ref. [5][6][7]). Instead, our approach with 34 instances already achieves optimized bandwidth consumption. This is very important as a lower number of instances lowers the orchestration overhead for network operators.

Further, the number of NFV nodes increases as the number of SC instances increases. Indeed, as SC mappings become more distributed, more nodes are being used for hosting virtual functions. In Fig. 3(a) 11 nodes are NFV enabled for 34 different SC mappings. For a network operator, capital expenditure in making 11 out of 14 nodes capable of hosting VNFs is very high. So, operators may want to minimize the number of NFV nodes while also trying to reduce bandwidth consumption by using multiple SC mapping instances. This led us to explore how the bandwidth consumption varies when the numbers of NFV nodes are limited.

Figure 3(b) shows the bandwidth consumption for SC mapping instances for various K values. When $K = 1$, all traffic flows have to traverse the one node in the network and hence, the number of instances does not affect bandwidth consumption. At $K = 2$, deploying more than 8 instances does not improve bandwidth utilization. For $K = 3$ and 16 instances we are able to achieve close to 10% of the minimum bandwidth utilization. Similarly, at $K = 4$ and 24 instances, we reach within 5% of the optimal bandwidth consumption. The bandwidth consumption comes to within 1% of the optimal when K=5 and 36 instances. Thus, we establish that we can achieve near-to-optimal bandwidth consumption by a using a relatively small number of instances and nodes.

Figures 3(c) and 3(d) corroborate our findings in Figs. 3(a) and 3(b). With the COST239 network too, we can achieve near-to-optimal bandwidth consumption with a small number of SC instances and nodes.

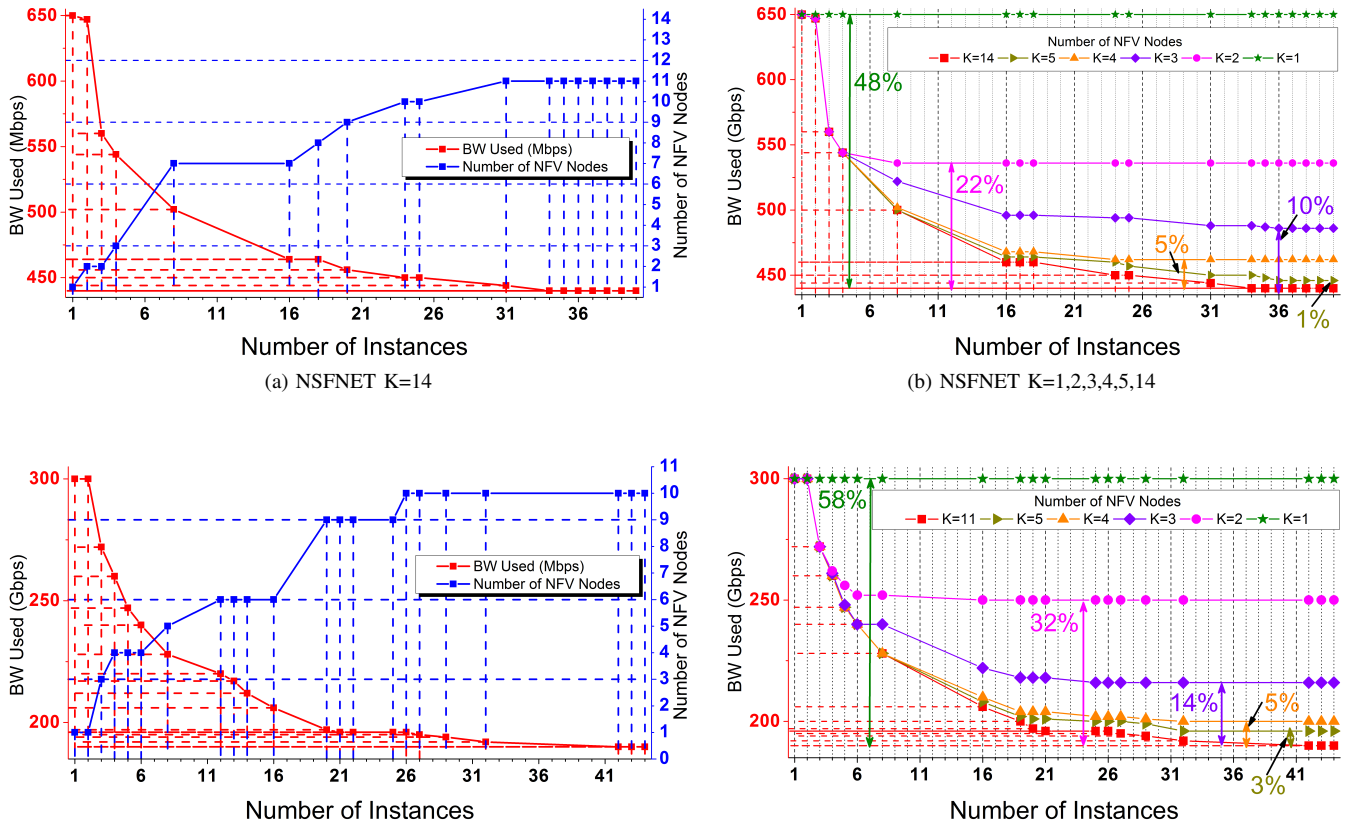

(a) NSFNET K=14

(b) NSFNET K=1,2,3,4,5,14

(c) COST239 K=11

(d) COST239 K=1,2,3,4,5,11

Fig. 3: Bandwidth vs. Number of NFV Nodes in NSFNET and COST239 networks.

VII. Conclusion

We introduce the problem of multiple service chain (SC) mapping with multiple SC instances in high traffic. We developed a column generation model along with a Shortest Path Traffic Grouping (SPTG) heuristic which results in a scalable linear model, thereby solving this complex problem in a relatively small amount of time. Further, we demonstrate that a near-to-optimal network resource consumption can be achieved with a relatively small number of SC instances and NFV nodes. This is critical in order to reduce the network operator's orchestration overhead and capital expenditures.

Acknowledgment

This work was supported by NSF Grant No. CNS-1217978.

References

[1] ETSI, "Network functions virtualisation: Introductory white paper," portal.etsi.org/NFV/NFV_White_Paper.pdf, 2012.

[2] IETF, "Network service chaining problem statement," https://tools.ietf.org/html/draft-quinn-nsc-problem-statement-00, 2013.

[3] S. Mehraghdam, M. Keller, and H. Karl, "Specifying and placing chains of virtual network functions," in *IEEE International Conference on Cloud Networking (CloudNet)*, 2014, pp. 7–13.

[4] H. Moens and F. De Turck, "VNF-P: A model for efficient placement of virtualized network functions," in *10th International Conference on Network and Service Management (CNSM)*, Nov. 2014, pp. 418–423.

[5] B. Addis, D. Belabed, M. Bouet, and S. Secci, "Virtual network functions placement and routing optimization," *https://hal.inria.fr/hal-01170042/*, 2015.

[6] M. C. Luizelli, L. R. Bays, L. S. Buriol, M. P. Barcellos, and L. P. Gaspary, "Piecing together the NFV provisioning puzzle: Efficient placement and chaining of virtual network functions," in *IFIP/IEEE IM*, May 2015, pp. 98–106.

[7] M. Bari, S. Chowdhury, R. Ahmed, and R. Boutaba, "On orchestrating virtual network functions in NFV," *Computing Research Repository*, vol. abs/1503.06377, 2015. [Online]. Available: http://arxiv.org/abs/1503.06377

[8] M. Savi, M. Tornatore, and G. Verticale, "Impact of processing costs on service chain placement in network functions virtualization," in *IEEE NFV-SDN*, 2015.

[9] N. Huin, B. Jaumard, and F. Giroire, "Optimization of Network Service Chain Provisioning." [Online]. Available: https://hal.inria.fr/hal-01476018

[10] A. Gupta, M. Habib, P. Chowdhury, M. Tornatore, and B. Mukherjee, "Joint Virtual Network Function Placement and Routing of Traffic in Operator Networks," *Technical Report, UC Davis*, 2015.

[11] B. Jaumard, C. Meyer, and B. Thiongane, "On column generation formulations for the RWA problem," *Discrete Applied Mathematics*, vol. 157, pp. 1291–1308, 2009.

[12] B. Jaumard and M. Daryalal, "Efficient spectrum utilization in large scale RWA problems," *IEEE/ACM Transactions on Networking*, pp. 1 – 16, 2017.

[13] A. Gupta, B. Jaumard, M. Tornatore, and B. Mukherjee, "Multiple Service Chain Placement and Routing in a Network-enabled Cloud," *CoRR*, vol. abs/1611.03197, 2016.

[14] M. F. Habib, M. Tornatore, M. De Leenheer, F. Dikbiyik, and B. Mukherjee, "Design of disaster-resilient optical datacenter networks," *Journal of Lightwave Technology*, vol. 30, no. 16, pp. 2563–2573, 2012.